High Availability

Palo Alto Networks®

PAN-OS® Administrator's Guide
Version 6.0

Copyright © 2007-2014 Palo Alto Networks

Contact Information

Corporate Headquarters:

Palo Alto Networks
4401 Great America Parkway
Santa Clara, CA 95054

http://www.paloaltonetworks.com/contact/contact/

About this Guide

This guide provides the concepts and solutions to help you get the most out of your Palo Alto Networks next-generation firewalls.

For more information, refer to the following sources:

- For start-to-finish instruction on how to set up a new firewall, refer to the Palo Alto Networks Getting Started Guide.
- For access to the complete technical documentation set, go to https://www.paloaltonetworks.com/documentation.
- For access to the knowledge base and discussion forums, go to https://live.paloaltonetworks.com.
- To contact support, for information on the support programs, or to manage your account or devices, go to https://support.paloaltonetworks.com.
- For the latest release notes, go to the Software Updates page at https://support.paloaltonetworks.com/Updates/SoftwareUpdates.

To provide feedback on the documentation, please write to us at:

documentation@paloaltonetworks.com

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2007-2015 Palo Alto Networks. All rights reserved.

Palo Alto Networks and PAN-OS are registered trademarks of Palo Alto Networks, Inc.

P/N 810-000211-00C 03.25.2014

Revision Date: March 23, 2015

High Availability

High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity.

The Palo Alto Networks firewalls support stateful active/passive or active/active high availability with session and configuration synchronization. Some models of the firewall, such as the VM-Series firewall and the PA-200, only support HA lite without session synchronization capability. The following topics provide more information about high availability and how to configure it in your environment.

- HA Overview
- HA Concepts
- Set Up Active/Passive HA
- HA Resources

Copyright © 2007-2014 Palo Alto Networks

HA Overview

On Palo Alto Networks firewalls, you can set up two devices as an HA pair. HA allows you to minimize downtime by making sure that an alternate device is available in the event that the primary device fails. The devices use dedicated or in-band HA ports on the firewall to synchronize data—network, object, and policy configurations—and to maintain state information. Device specific configuration such as management port IP address or administrator profiles, HA specific configuration, log data, and the Application Command Center (ACC) information is not shared between devices. For a consolidated application and log view across the HA pair, you must use Panorama, the Palo Alto Networks centralized management system.

When a failure occurs on the active device and the passive device takes over the task of securing traffic, the event is called a failover. The conditions that trigger a failover are:

- One or more of the monitored interfaces fail. (Link Monitoring)
- One or more of the destinations specified on the device cannot be reached. (Path Monitoring)
- The device does not respond to heartbeat polls. (Heartbeat Polling and Hello messages)

After you understand the HA Concepts, continue to Set Up Active/Passive HA.

Copyright © 2007-2014 Palo Alto Networks

HA Concepts

The following topics provide conceptual information about how HA works on a Palo Alto Networks firewall:

- HA Modes
- HA Links and Backup Links
- Device Priority and Preemption
- Failover Triggers
- HA Timers

HA Modes

You can set up the firewalls for HA in two modes:

- **Active/Passive**— One device actively manages traffic while the other is synchronized and ready to transition to the active state, should a failure occur. In this configuration, both devices share the same configuration settings, and one actively manages traffic until a path, link, system, or network failure occurs. When the active device fails, the passive device takes over seamlessly and enforces the same policies to maintain network security. Active/passive HA is supported in the virtual wire, Layer 2 and Layer 3 deployments. For information on setting up your devices in an active/passive configuration, see Configure Active/Passive HA.

The PA-200 and the VM-Series firewalls support a lite version of active/passive HA. HA lite provides configuration synchronization and some runtime data synchronization such as IPSec security associations. It does not support any session synchronization, and therefore, HA Lite does not offer stateful failover.

- **Active/Active**— Both the devices in the pair are active and processing traffic, and work synchronously to handle session setup and session ownership. The active/active deployment is supported in virtual wire and Layer 3 deployments, and is only recommended for networks with asymmetric routing. For information on setting up the devices in an active/active configuration, refer to the Active/Active High Availability Tech Note.

HA Links and Backup Links

The devices in an HA pair use HA links to synchronize data and maintain state information. Some models of the firewall have dedicated HA ports—Control link (HA1) and Data link (HA2), while others require you to use the in-band ports as HA links.

On devices with dedicated HA ports such as the PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050 firewalls (see HA Ports on the PA-7050 Firewall), use the dedicated HA ports to manage communication and synchronization between the devices. For devices without dedicated HA ports such as the PA-200, PA-500, and PA-2000 Series firewalls, as a best practice use the management port for the HA1 link to allow for a direct connection between the management planes on the devices, and an in-band port for the HA2 link.

Copyright © 2007-2014 Palo Alto Networks

The HA1 and HA2 links provide synchronization for functions that reside on the management plane. Using the dedicated HA interfaces on the management plane is more efficient than using the in-band ports as this eliminates the need to pass the synchronization packets over the dataplane.

- **Control Link**: The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing, and User-ID information. This link is also used to synchronize configuration changes on either the active or passive device with its peer. The HA1 link is a Layer 3 link and requires an IP address.
  Ports used for HA1: TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP).

- **Data Link**: The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between devices in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active device to the passive device. The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.
  Ports used for HA2: The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.
  Additionally, an HA3 link is used in Active/Active HA deployments. When there is an asymmetric route, the HA3 link is used for forwarding packets to the HA peer that owns the session. The HA3 link is a Layer 2 link and it does not support Layer 3 addressing or encryption.

- **Backup Links**: Provide redundancy for the HA1 and the HA2 links. In-band ports are used as backup links for both HA1 and HA2. Consider the following guidelines when configuring backup HA links:
  - The IP addresses of the primary and backup HA links must not overlap each other.
  - HA backup links must be on a different subnet than the primary HA links.
  - HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses port 28770 and 28260.

Palo Alto Networks recommends enabling heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.

HA Ports on the PA-7050 Firewall

For HA connectivity on the PA-7050, refer to the following table for details on which ports on the Switch Management Card (SMC) are mandated and where ports on the Network Processing Card (NPC) are suitable. For an overview of the Modules and Interface cards on the PA-7050 firewall, refer to the PA-7050 Hardware Reference Guide.

The following ports on the SMC are designed for HA connectivity:

Copyright © 2007-2014 Palo Alto Networks

HA Links and Backup Links	Ports on the SMC	Description
Control Link	HA1-A Speed: Ethernet 10/100/1000	Used for HA control and synchronization. Connect this port directly from the HA1-A port on the first device to the HA1-A on the second device in the pair, or connect them together through a switch or router. HA1 cannot be configured on NPC data ports or the MGT port.
Control Link Backup	HA1-B Speed: Ethernet 10/100/1000 port	Used for HA control and synchronization as a backup for HA1-A. Connect this port directly from the HA1-B port on the first device to the HA1-B on the second device in the pair, or connect them together through a switch or router. HA1 Backup cannot be configured on NPC data ports or the MGT port.
Data Link	HSCI-A (High Speed Chassis Interconnect)	Quad Port SFP (QSFP) interfaces used to connect two PA-7050 firewalls in an HA configuration. Each port is comprised of four 10 gigabit links internally for a combined speed of 40 gigabits and is used for HA2 data link in an active/passive configuration. When in active/active mode, the port is also used for HA3 packet forwarding for asymmetrically routed sessions that require Layer 7 inspection for App-ID and Content-ID. In a typical installation, HSCI-A on the first chassis connects directly to HSCI-A on the second chassis and HSCI-B on the first chassis connects to HSCI-B on the second chassis. This will provide full 80 gigabit transfer rates. In software, both ports (HSCI-A and HSCI-B) are treated as one HA interface. The HSCI ports are not routable and must be connected directly to each other. Palo Alto Networks recommends using the dedicated HSCI ports for both HA2 and HA3 connections. However, the HA2 and HA3 links can be configured on NPC data ports, if needed.
Data Link Backup	HSCI-B (High Speed Chassis Interconnect)	The Quad Port SFP (QSFP) interfaces (see description above) in the HSCI-B port is used to increase the bandwidth for HA2/HA3 purposes. The HSCI ports are not routable and must be connected directly to each other. Palo Alto Networks recommends using the dedicated HSCI-B ports for both HA2 and HA3 backup connections. The HA2/HA3 backup link can be configured on the NPC data ports, if needed.

Device Priority and Preemption

The devices in an HA pair can be assigned a *device priority* value to indicate a preference for which device should assume the active role and manage traffic. If you need to use a specific device in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each

Copyright © 2007-2014 Palo Alto Networks

device. The device with the lower numerical value, and therefore *higher priority*, is designated as active and manages all traffic on the network. The other device is in a passive state, and synchronizes configuration and state information with the active device so that it is ready to transition to an active state should a failure occur.

By default, preemption is disabled on the firewalls and must be enabled on both devices. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical value) to resume as active after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

Failover Triggers

When a failure occurs on the active device and the passive device takes over the task of securing traffic, the event is called a failover. A failover is triggered when a monitored metric on the active device fails. The metrics that are monitored for detecting a device failure are:

- **Heartbeat Polling and Hello messages**

  The firewalls use hello message and heartbeats to verify that the peer device is responsive and operational. Hello messages are sent from one peer to the other at the configured *Hello Interval* to verify the state of the device. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the devices are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds. For details on the HA timers that trigger a failover, see HA Timers.

- **Link Monitoring**

  The physical interfaces to be monitored are grouped into a link group and their state (link up or link down) is monitored. A link group can contain one or more physical interfaces. A device failure is triggered when any or all of the interfaces in the group fail. The default behavior is failure of any one link in the link group will cause the device to change the HA state to non-functional to indicate a failure of a monitored object.

- **Path Monitoring**

  Monitors the full path through the network to mission-critical IP addresses. ICMP pings are used to verify reachability of the IP address. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail, and a device failure is triggered when any or all of the IP addresses monitored become unreachable. The default behavior is any one of the IP addresses becoming unreachable will cause the device to change the HA state to non-functional to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator places the device is a suspended state or if preemption occurs.

On the PA-3000 Series, PA-5000 Series, and PA-7050 firewalls, a failover can occur when an internal health check fails. This health check is not configurable and is enabled to verify the operational status for all the components within the firewall.

Copyright © 2007-2014 Palo Alto Networks

HA Timers

High Availability (HA) timers are used to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, you can select from three profiles have been added: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

The following table describes each timer included in the profiles and the current preset values across the different hardware models; these values are for current reference only and can change in a subsequent release.

Recommended/Aggressive HA Timer Values by Platform

Timers	Description	PA-7050 PA-5000 Series PA-4000 Series PA-3000 Series	PA-2000 Series PA-500 Series PA-200 Series VM-Series	Panorama VM M-100
Monitor fail hold up time	The interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices.	0/0	0/0	0/0
Preemption hold time	Time a passive or active-secondary device will wait before taking over as the active or active-primary device.	1/1	1/1	1/1
Heartbeat interval	The frequency at which the HA peers exchange heartbeat messages in the form of an ICMP ping.	1000/1000	2000/1000	2000/1000
Promotion hold time	Time that the passive device (in active/passive mode) or the active-secondary device (in active/active mode) will wait before taking over as the active or active-primary device after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.	2000/500	2000/500	2000/500

Copyright © 2007-2014 Palo Alto Networks

Timers	Description	PA-7050 PA-5000 Series PA-4000 Series PA-3000 Series	PA-2000 Series PA-500 Series PA-200 Series VM-Series	Panorama VM M-100
Additional master hold up time	This time interval is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active device in active/passive mode and to the active-primary device in active/active mode. This timer is recommended to avoid a failover when both devices experience the same link/path monitor failure simultaneously.	500/500	500/500	7000/5000
Hello interval	The time interval in milliseconds between the hello packets that are sent to verify that the HA functionality on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms.	8000/8000	8000/8000	8000/8000
Maximum no. of flaps	A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. This value indicates the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16, default 3).	3/3	3/3	Not Applicable

Copyright © 2007-2014 Palo Alto Networks

Set Up Active/Passive HA

- Prerequisites for Active/Passive HA
- Configuration Guidelines for Active/Passive HA
- Configure Active/Passive HA
- Define HA Failover Conditions
- Verify Failover

Copyright © 2007-2014 Palo Alto Networks

Prerequisites for Active/Passive HA

To set up high availability on your Palo Alto Networks firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—both the devices in the pair must be of the same hardware model or virtual machine model.
- **The same PAN-OS version**—both the devices should be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases. They must also both have the same multiple virtual systems capability (single or multi vsys).
- **The same type of interfaces**—dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - Determine the IP address for the HA1 (control) connection between the device pair. The HA1 IP address for both peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For devices without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both devices. However, because the management ports will not be directly cabled between the devices, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 *only* if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
- **The same set of licenses**—Licenses are unique to each device and cannot be shared between the devices. Therefore, you must license both devices identically. If both devices do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.

If you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration, it is recommended that you perform a factory reset on the new firewall. This will ensure that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary device to the newly introduced device with the clean config.

Copyright © 2007-2014 Palo Alto Networks

Configuration Guidelines for Active/Passive HA

To set up an active (PeerA) passive (PeerB) pair in HA, you must configure some options identically on both devices and some independently (non-matching) on each device. These HA settings are not synchronized between the devices. For details on what is/is not synchronized, refer to HA Synchronization.

To proceed with the instructions on configuring the devices in HA, see Configure Active/Passive HA.

The following table lists the settings that you must configure identically on both devices:

Identical Configuration Settings on PeerA and PeerB
• HA must be enabled on both devices. • Both device must have the same Group ID value. The Group ID value is used to create a virtual MAC address for all the configured interfaces. The format of the virtual MAC is 00-1B-17:00: xx: yy where 00-1B-17: vendor ID; 00: fixed; xx: HA group ID; yy: interface ID. When a new active device takes over, Gratuitous ARPs are sent from each of the connected interfaces of the new active member to inform the connected Layer 2 switches of the virtual MAC address' new location. • If using in-band ports, the interfaces for the HA1 and HA2 links must be set to type HA. • The HA mode must be set to Active Passive. • If required, preemption must be enabled on both devices. The device priority value, however, must not be identical. • If required, encryption on the HA1 link (for communication between the HA peers) must be configured on both devices. • Based on the combination of HA1 and HA1 Backup ports you are using, use the following recommendations to decide whether you should enable heartbeat backup: • HA1: Dedicated HA1 port HA1 Backup: In-band port Recommendation: Enable Heartbeat Backup • HA1: Dedicated HA1 port HA1 Backup: Management port Recommendation: Do not enable Heartbeat Backup • HA1: In-band port HA1 Backup: In-band port Recommendation: Enable Heartbeat Backup • HA1: Management port HA1 Backup: In-band port Recommendation: Do not enable Heartbeat Backup

The following table lists the settings that must be configured independently on each device:

Copyright © 2007-2014 Palo Alto Networks

Independent Configuration Settings	PeerA	PeerB
Control Link	IP address of the HA1 link configured on this device (PeerA).	IP address of the HA1 link configured on this device (PeerB).
	For devices without dedicated HA ports, use the management port IP address for the control link.	
Data Link The data link information is synchronized between the devices after HA is enabled and the control link is established between the devices.	By default, the HA2 link uses Ethernet/Layer 2. If using a Layer 3 connection, configure the IP address for the data link on this device (PeerA).	By default, the HA2 link uses Ethernet/Layer 2. If using a Layer 3 connection, configure the IP address for the data link on this device (PeerB).
Device Priority (required, if preemption is enabled)	The device you plan to make active must have a lower numerical value than its peer. So, if Peer A is to function as the active device, keep the default value of 100 and increment the value on PeerB.	If PeerB is passive, set the device priority value to a number larger than that on PeerA. For example, set the value to 110.
Link Monitoring—Monitor one or more physical interfaces that handle vital traffic on this device and define the failure condition.	Select the physical interfaces on the firewall that you would like to monitor and define the failure condition (all or any) to trigger a failover.	Pick a similar set of physical interfaces that you would like to monitor on this firewall and define the failure condition (all or any) to trigger a failover.
Path Monitoring—Monitor one or more destination IP addresses that the firewall can use ICMP pings to ascertain responsiveness.	Define the failure condition (all or any), ping interval and the ping count. This is particularly useful for monitoring the availability of other interconnected networking devices. For example, monitor the availability of a router that connects to a server, connectivity to the server itself, or some other vital device that is in the flow of traffic. Make sure that the node/device that you are monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a a path monitoring failure and trigger a failover.	Pick a similar set of devices or destination IP addresses that can be monitored for determining the failover trigger for PeerB. Define the failure condition (all or any), ping interval and the ping count.

Copyright © 2007-2014 Palo Alto Networks

Configure Active/Passive HA

The following procedure shows how to configure a pair of firewalls in an active/passive deployment as depicted in the following example topology.

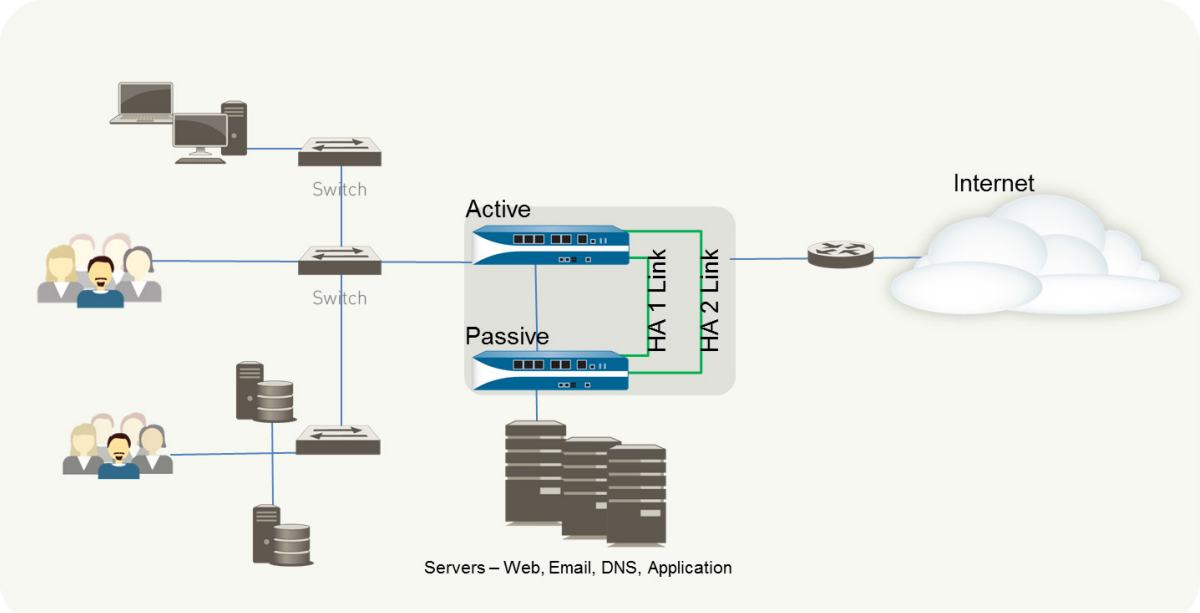

Connect and Configure the Devices	
Step 1 Connect the HA ports to set up a physical connection between the devices.	• For devices with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on the device pair. Use a crossover cable if the devices are directly connected to each other. • For devices without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both devices. Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network.
Pick a device in the pair and complete these tasks:	
Step 2 Enable ping on the management port. Enabling ping allows the management port to exchange heartbeat backup information.	1. Select **Device > Setup > Management** and then click the Edit icon in the Management Interface Settings section of the screen. 2. Select **Ping** as a service that is permitted on the interface.
Step 3 If the device does not have dedicated HA ports, set up the data ports to function as HA ports. For devices with dedicated HA ports continue to Step 4.	1. Select **Network > Interfaces.** 2. Confirm that the link is up on the ports that you want to use. 3. Select the interface and set the interface type to HA. 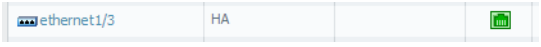 4. Set the **Link Speed** and **Link Duplex** settings, as appropriate.

Copyright © 2007-2014 Palo Alto Networks

Connect and Configure the Devices (Continued)	
Step 4 Set up the control link connection. This example shows an in-band port that is set to interface type HA. For devices that use the management port as the control link, the IP address information is automatically pre-populated.	1. In **Device > High Availability > General**, edit the **Control Link (HA1)** section. 2. Select the interface that you have cabled for use as the HA1 link in the **Port** drop down menu. Set the IP address and netmask. Enter a Gateway IP address only if the HA1 interfaces are on separate subnets. Do not add a gateway if the devices are directly connected. 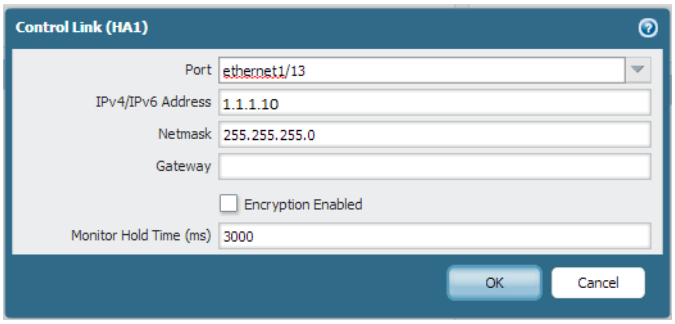
Step 5 (Optional) Enable encryption for the control link connection. This is typically used to secure the link if the two devices are not directly connected, that is if the ports are connected to a switch or a router.	1. Export the HA key from a device and import it into the peer device. a. Select **Device > Certificate Management > Certificates.** b. Select **Export HA key.** Save the HA key to a network location that the peer device can access. c. On the peer device, navigate to **Device > Certificate Management > Certificates,** and select **Import HA key** to browse to the location that you saved the key and import it in to the peer device. 2. Select **Device > High Availability > General**, edit the **Control Link (HA1)** section. 3. Select **Encryption Enabled.**
Step 6 Set up the backup control link connection.	1. In **Device > High Availability > General**, edit the **Control Link (HA1 Backup)** section. 2. Select the HA1 backup interface and set the IP address and netmask. 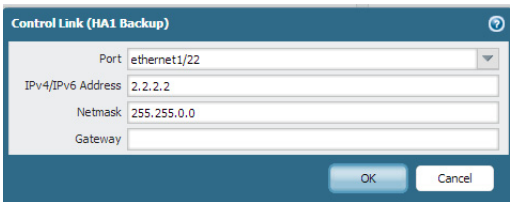

Copyright © 2007-2014 Palo Alto Networks

Connect and Configure the Devices (Continued)	
Step 7 Set up the data link connection (HA2) and the backup HA2 connection between the devices.	1. In **Device > High Availability > General**, edit the **Data Link (HA2)** section. 2. Select the interface for the data link connection. 3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode. 4. If you use IP or UDP as the transport method, enter the IP address and netmask. 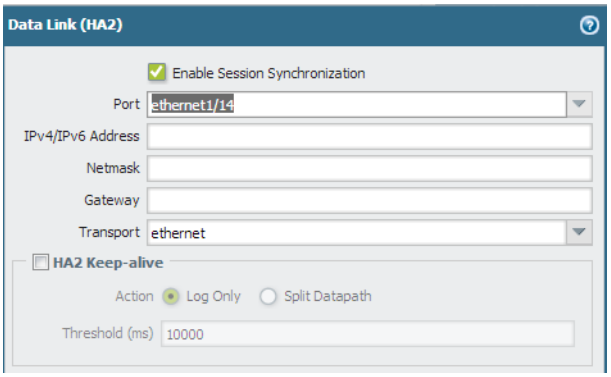 5. Verify that **Enable Session Synchronization** is selected. 6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. For active/passive configuration, a critical system log message is generated when an HA2 keep-alive failure occurs. You can configure the HA2 keep-alive option on both devices, or just one device in the HA pair. If the option is only enabled on one device, only that device will send the keep-alive messages. The other device will be notified if a failure occurs. 7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the IP address and netmask.
Step 8 Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port. You do not need to enable heartbeat backup if you are using the management port for the control link.	1. In **Device > High Availability > General**, edit the **Election Settings** section. 2. Select **Heartbeat Backup**. The heartbeat backup link is used for transmitting redundant heartbeats and hello messages. To allow the heartbeats to be transmitted between the devices, you must verify that the management port across both peers can route to each other.

Copyright © 2007-2014 Palo Alto Networks

Connect and Configure the Devices (Continued)	
Step 9 Set the device priority and enable preemption. This setting is only required if you wish to make sure that a specific device is the preferred active device. For information, see Device Priority and Preemption.	1. In **Device > High Availability > General**, edit the **Election Settings** section. 2. Set the numerical value in **Device Priority**. Make sure to set a lower numerical value on the device that you want to assign a higher priority to. If both firewalls have the same device priority value, the firewall with the lowest MAC address on the HA1 control link will become the active device. 3. Select **Preemptive**. You must enable preemptive on both the active and the passive device.
Step 10 (Optional) Modify the failover timers. By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments.	1. In **Device > High Availability > General**, edit the **Election Settings** section. 2. Select the **Aggressive** profile for triggering failover faster; select **Advanced** to define custom values for triggering failover in your set up. To view the preset value for an individual timer included in a profile, select **Advanced** and click **Load Recommended** or **Load Aggressive**. The preset values for your hardware model will be displayed on screen.
Step 11 (Optional, only configured on the passive device) Modify the link status of the HA ports on the passive device. The passive link state is **shutdown**, by default. After you enable HA, the link state for the HA ports on the active device will be green and those on the passive device will be down and display as red.	Setting the link state to **Auto** allows for reducing the amount of time it takes for the passive device to take over when a failover occurs and it allows you to monitor the link state. To enable the link status on the passive device to stay up and reflect the cabling status on the physical interface: 1. In **Device > High Availability > General**, edit the **Active Passive Settings** section. 2. Set the **Passive Link State** to **Auto**. The auto option decreases the amount of time it takes for the passive device to take over when a failover occurs. Although the interface displays green (as cabled and up) it continues to discard all traffic until a failover is triggered. When you modify the passive link state, make sure that the adjacent devices do not forward traffic to the passive firewall based only on the link status of the device.

Copyright © 2007-2014 Palo Alto Networks

Connect and Configure the Devices (Continued)	
Step 12 Enable HA.	1. Select **Device > High Availability > General,** edit the **Setup** section. 2. Select **Enable HA.** 3. Set a **Group ID**. This ID uniquely identifies each HA pair on your network, and is essential if you have multiple HA pairs that share the same broadcast domain on your network. 4. Set the mode to **Active Passive**. 5. Select **Enable Config Sync.** This setting enables the synchronization of the configuration settings between the active and the passive device. 6. Enter the IP address assigned to the control link of the peer device in **Peer HA1 IP Address**. 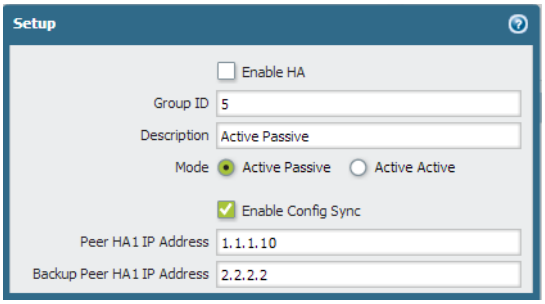 For devices without dedicated HA ports, if the peer uses the management port for the HA1 link, enter the management port IP address of the peer. 7. Enter the **Backup HA1 IP Address**.
Step 13 Save your configuration changes.	Click **Commit**.
Step 14 Complete Step 2 through Step 13 on the other device in the HA pair.	
Step 15 After you finish configuring both devices, verify that the devices are paired in active/passive HA.	1. Access the **Dashboard** on both devices, and view the **High Availability** widget. 2. On the active device, click the **Sync to peer** link. 3. Confirm that the devices are paired and synced, as shown below:

Copyright © 2007-2014 Palo Alto Networks

Connect and Configure the Devices (Continued)

On the passive device: The state of the local device should display **passive** and the configuration is **synchronized.**

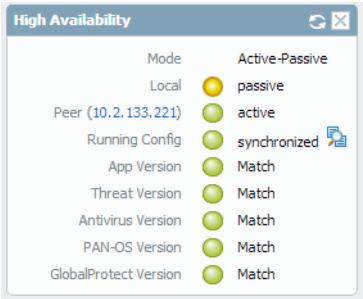

On the active device: The state of the local device should display **active** and the configuration is **synchronized**.

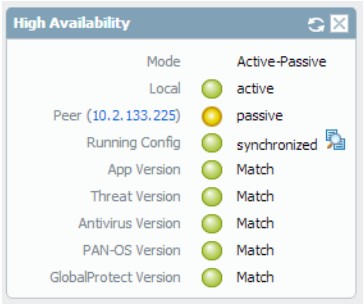

Copyright © 2007-2014 Palo Alto Networks

Define HA Failover Conditions

Configure the Failover Triggers		
Step 1	To configure link monitoring, define the interfaces that you would like to monitor. A change in the link state of these interface will trigger a failover.	1. Select **Device > High Availability > Link and Path Monitoring.** 2. In the **Link Group** section, click **Add.** 3. Name the **Link Group, Add** the interfaces to monitor, and select the **Failure Condition** for the group. The Link group you define is added to the **Link Group** section.
Step 2	(Optional) Modify the failure condition for the Link Groups that you configured (in the preceding step) on the device. By default, the device will trigger a failover when any monitored link fails.	1. Select the **Link Monitoring** section. 2. Set the **Failure Condition** to **All.** The default setting is **Any**.
Step 3	To configure path monitoring, define the destination IP addresses that the firewall should ping to verify network connectivity.	1. In the **Path Group** section of the **Device > High Availability > Link and Path Monitoring** tab, pick the **Add option for your set up:** Virtual Wire, VLAN, or Virtual Router. 2. Select the appropriate item from the drop-down list for the **Name** and **Add** the IP addresses (source and/or destination, as prompted) that you wish to monitor. Then select the **Failure Condition** for the group. The path group you define is added to the **Path Group** section. 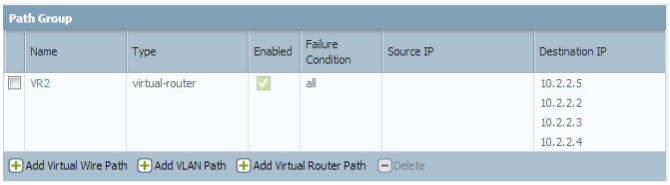
Step 4	(Optional) Modify the failure condition for all Path Groups configured on the device. By default, the device will trigger a failover when any monitored path fails.	Set the **Failure Condition** to **All.** The default setting is **Any**.
Step 5	Save your changes.	Click **Commit**.

If you are using SNMPv3 to monitor the firewalls, note that the SNMPv3 Engine ID is unique to each device; the EngineID is not synchronized between the HA pair and therefore, allows you to independently monitor each device in the HA pair. For information on setting up SNMP, see Set Up SNMP Trap Destinations.

Because the EngineID is generated using the device's unique serial number, on the VM-Series firewall you must apply a valid license in order to obtain a unique EngineID for each firewall.

Copyright © 2007-2014 Palo Alto Networks

Verify Failover

To test that your HA configuration works properly trigger a manual failover and verify that the devices transition states successfully.

Verify Failover	
Step 1 Suspend the active device.	Click the **Suspend local device** link on the **Device > High Availability > Operational Commands** tab. 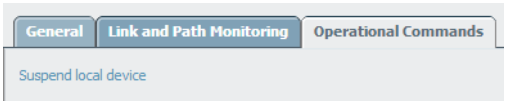
Step 2 Verify that the passive device has taken over as active.	On the **Dashboard**, verify that the state of the passive device changes to **active** in the **High Availability** widget. 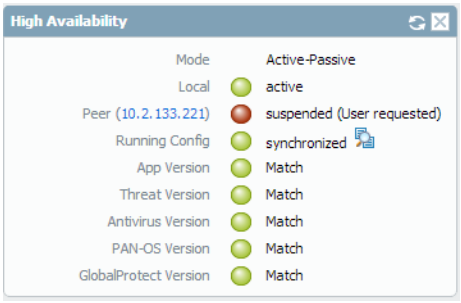
Step 3 Restore the suspended device to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled.	1. On the device you previously suspended, select the **Make local device functional** link on the **Device > High Availability > Operational Commands** tab. 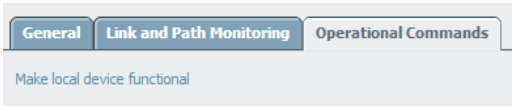 2. In the **High Availability** widget on the **Dashboard,** confirm that the device has taken over as the active device and that the peer is now in a passive state. 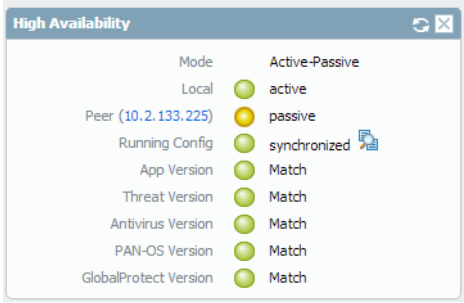

Copyright © 2007-2014 Palo Alto Networks

HA Resources

For more information on HA, refer to the following sources:

- Active/Active HA
- High Availability Synchronization
- High Availability Failover Optimization
- Upgrading an HA pair
- Examples: Deploying HA

Copyright © 2007-2014 Palo Alto Networks

Copyright © 2007-2014 Palo Alto Networks